Attorney Docket No.: 20510-19.00US

PATENT APPLICATION

APPARATUS AND METHOD OF LOOP AND RATE DEPENDENT POWER CUTBACK

Inventors: Alberto Ginesi, a citizen of Canada, residing at
111 Covington Place, Nepean, Ontario K2G 6B6 Canada

Francois Tremblay, a citizen of Canada, residing at
103 St. Therese, Hull, Quebec J9A 2M5 Canada

Scott McClennon, a citizen of Canada, residing at
209 Chandler Avenue, Ottawa, Ontario K2C 0G3 Canada

Assignee: Catena Networks, Inc.
303 Twin Dolphin Drive, Suite 600
Redwood Shores, CA 94065

Entity: Small business concern

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834
Tel: 650-326-2400

APPARATUS AND METHOD OF LOOP AND RATE DEPENDENT POWER CUTBACK

CROSS-REFERENCES TO RELATED APPLICATIONS

[01] NOT APPLICABLE

STATEMENT AS TO RIGHTS TO INVENTIONS MADE UNDER FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

[02] NOT APPLICABLE

REFERENCE TO A "SEQUENCE LISTING," A TABLE, OR A COMPUTER PROGRAM LISTING APPENDIX SUBMITTED ON A COMPACT DISK.

[03] NOT APPLICABLE

BACKGROUND OF THE INVENTION

[04] The present invention relates generally to transfer of data using Digital Subscriber Loop (DSL) technology, and specifically to a method for reducing transmitter power used for the transfer.

[05] Remote access and retrieval of data is becoming increasingly popular in data communication. The proliferation of the Internet has provided a vast network of information that is available to the general public. As the Internet grows and technology advances, this information is becoming increasingly voluminous and the details are become increasingly intricate. What used to comprise mainly text information has grown to include still and moving images as well as sound. The increase in the volume of information to be transferred has presented a need for a high-speed Internet connection, since traditional telephone modems communicate at speeds too slow for efficient communication.

[06] One proposal for high-speed communication is the introduction of Digital Subscriber Line (DSL) technology. The various DSL technologies include asymmetric DSL (ADSL), high-speed DSL (HDSL), symmetric DSL (SDSL), Symmetric High-Bit-Rate DSL (SHDSL) and Integrated Services Digital Network (ISDN) Basic Rate Interface (BRI) DSL systems. One of the most attractive features of DSL is that it is implemented using an infrastructure that already exists. DSL shares copper twisted pair lines

typically used for telephone communication. However, only a small portion of the available bandwidth of the twisted pair line (0 to 4kHz) is used for Plain Old Telephone Service (POTS). DSL takes advantage of the available frequency spectrum from 4kHz to approximately 1.1MHz for transmitting data.

[07] ADSL is currently the most practical form of DSL technology, and therefore the most widely implemented. ADSL is asymmetric in that its downstream (DS or D/S -- to a subscriber) capacity is larger than its upstream (US or U/S -- from the subscriber) capacity. An ADSL transceiver unit at a central office a central office or remote loop carrier (ATU-C) is used for sending downstream information and receiving upstream information. An ADSL transceiver unit at a remote location or user end (ATU-R) is used for receiving downstream information and sending upstream information. Typically, a Discrete Multi-tone (DMT) scheme is used. The spectrum from 4kHz to 1.1MHz is divided into 256 sub-channels, or tones, each having a bandwidth of 4.3125kHz. Each sub-channel uses Quadrature Amplitude Modulation (QAM) to carry 2 to 15 bits/QAM symbol.

[08] A predefined power level is used by the ATU-C for transmitting the downstream signal. At the ATU-R, the achievable downstream rate is a function of the received signal level and of the receiver input-referenced noise levels. Where the loop attenuation is modest and the target downstream rate is lower than the achievable rate, the predefined ATU-C transmission level can be much higher than necessary. Therefore, an unnecessary amount of power is consumed by the ATU-C and additional crosstalk noise is induced in adjacent DSL lines. Reducing the ATU-C transmit level would both save power at the ATU-C and reduce crosstalk noise, improving the quality of signal on adjacent lines. From a crosstalk perspective, it is beneficial if the ATU-C transmit power can be reduced during both initialization and steady state (also referred to as "Showtime") operation. If the transmit power and resulting crosstalk is only reduced on entry to Showtime, the higher crosstalk levels during initialization may cause excessive errors and even force re-initialization on adjacent DSL lines.

[09] The benefit of such a power cutback can be significant in Digital Loop Carrier (DLC) applications, for example, where subscriber loops are typically shorter than seen at central office (CO)-resident line interfaces and where power and thermal budgets are tight. Assuming a DMT transmit signal with peak-to-average-ratio (PAR) of 14.5dB, downstream power cutback can yield significant savings on line driver power alone.

[10] American National Standards Institute (ANSI) and International Telecommunications Union (ITU) compliant ATU-Cs reduce their transmitter power on very

short loops to avoid overloading the ATU-R. This reduction in downstream transmission power is often referred to as a "politeness cutback". The politeness cutback is determined in accordance with a loop loss estimate from an upstream signal path and applies on loops shorter than approximately 2-3kft 26AWG-equivalent. Other downstream transmitter power cutbacks are also specified in the ITU splitterless ADSL standard (G.992.2, which is also known as G.lite) for reducing the downstream signal level to the ATU-R. This cutback is designed to address ADSL signal levels that could cause distortion in the presence of an off-hook telephone set.

[11] However, even with the implementation of the politeness cutback and off-hook cutback, the ATU-C often transmits at a greater power than necessary and there is no agreed-upon mechanism to implement a general downstream power cutback. It is an object of the present invention to obviate or mitigate at least some of the above-mentioned disadvantages.

BRIEF SUMMARY OF THE INVENTION

[12] In accordance with an aspect of the present invention, there is provided a method for reducing power required for transmitting a signal from a first transceiver to a second transceiver. The method comprises the steps of estimating an excess amount of power used by the first transceiver for transmitting the signal, reducing the first transceiver's power use by the excess amount of power to a reduced power level, and transmitting the signal from the first transceiver using the reduced power level. The reduced power level achieves a transmission rate of the signal within a predefined tolerance of its preferable rate.

BRIEF DESCRIPTION OF THE DRAWINGS

[13] Embodiments of the invention will now be described by way of example only with reference to the following drawings in which:

[14] FIG. 1a is a graph illustrating a relationship between transmission rate and loop length for signals transmitted at a plurality of power levels, in the presence of forty-nine G.Lite disturbers;

[15] FIG. 1b is a graph illustrating a relationship between transmission rate and loop length for signals transmitted at a plurality of power levels, in the presence of ten HDSL disturbers;

[16] FIG. 1c is a graph illustrating a relationship between transmission rate and loop length for signals transmitted at a plurality of power levels, in the presence of five adjacent binder T1 disturbers;

[17] FIGS. 2a and 2b are flowcharts of the process steps according to embodiments of the present invention; and

[18] FIG. 3 is a block diagram of an ATU-C and an ATU-R that implement embodiments of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[19] ATU-C power requirements can be reduced when there is an excess signal-to-noise ratio (SNR) margin at the ATU-R receiver. It is possible to determine an amount of cutback possible given a target or maximum downstream data rate and a target or maximum downstream SNR margin. The target data rate is the rate at which the ATU-C is to transmit the downstream signal, and the SNR margin is a margin against SNR degradation for a specified bit-error-rate.

[20] Referring to FIGS. 1a through 1c, the graphs illustrate attainable downstream rates over loop lengths of 0-12kft (26AWG). Each graph has a plot illustrating the estimated transmission rates with respect to the loop length with a DS power cutback of 0dB, 6dB, 12dB, and 15dB in the presence of crosstalk from other DSLs on adjacent pairs. Full-rate ADSL (G.dmt and T1.413) are also illustrated for comparison. The attainable rates have been estimated assuming a 4dB SNR margin, a 3dB coding gain, use of downstream carriers 36-127, and a downstream receiver noise floor of -136.8dBm/Hz. The politeness cutback (of 0 to 12dB) is not shown in the graphs. However, the politeness cutback is only applied for loop lengths up to approximately 2kft (26AWG).

[21] Referring to FIG. 1a, the crosstalk scenario is caused by forty-nine (49) other G.lite disturbers. This case is typical of a residential neighborhood, where T1 or HDSL services in the same or adjacent binders is less common. It should be noted that if all local subscribers are being served off power cutback-capable digital loop carriers (DLCs), the G.lite crosstalk levels would actually be lower and the rates would be improved over those shown.

[22] Referring to FIG. 1b, the crosstalk scenario is caused by ten (10) HDSL disturbers. For this case, full capacity (1.5Mbps) is possible with 15dB of cutback for loop lengths of up to 9kft (26AWG), within Carrier Serving Area (CSA) loop engineering rules.

[23] Referring to FIG. 1c, the crosstalk scenario is caused by five (5) adjacent binder T1 disturbers. With T1 disturbers in adjacent binders, full capacity is possible on a loop of length of up to 9kft with approximately 6dB of cutback.

[24] In order to implement an appropriate power cutback, it is necessary to either obtain an estimate of the SNR at the receiver or develop a method for determining the amount of signal cutback possible while achieving the target downstream rate and SNR margin. A signal-to-additive-noise estimate (as opposed to signal-dependent-noise estimate, such as mis-equalization error) permits a more accurate estimate of the SNR degradation resulting from a signal level reduction. Once the cutback is determined, it is communicated between the ATU-C and the ATU-R and implemented by the ATU-C.

[25] Estimating the SNR and determining an optimal amount of downstream transmit signal power cutback is most easily accomplished by the ATU-R receiver. Thus far, some proposals have been made in the ITU standards body to permit this and/or require the ATU-R to minimize downstream transmit power via the per-DMT-carrier gains when the maximum downstream SNR margin is exceeded. These methods are not required or supported in the current ADSL standards nor do they address ATU-C transmit (and generated crosstalk) levels during initialization. As a result, it is preferable to have methods for estimating downstream SNR and/or potential cutback in the ATU-C.

[26] The techniques described herein describe a downstream power cutback mechanism that reduces downstream transmit power based on a metric of excess SNR or capacity at the downstream receiver (ATU-R). The following techniques can be implemented with an ATU-R compliant with the current ADSL standards (ITU-T G.992.1, G.992.2 and ANSI T1.413 issue 2).

[27] The amount of downstream transmit power reduction that can be tolerated for a given capacity target is a function of the loop and the crosstalk environment, which together determine the SNR per receiver sub-carrier. Unfortunately, this information is not available to the ATU-C. This information may be inferred from the bits per carrier *i* (B_is) and gains per carrier *i* (G_is -- also called power per carrier) sent during initialization. Specifically, bits and gains information from the remote-end is transmitted during R-B&G, toward the end of the initialization. However, at this point it is too late to implement a downstream power cutback without impacting the ATU-R, which will at least need to adapt its receiver automatic gain control (AGC).

[28] Therefore, in the absence of explicit information from the ATU-R, several approaches to initialization are disclosed that the potential for power savings without

significant capacity losses. For all cases, the amount of cutback versus the measured upstream power can be an operator-specified parameter, as will be appreciated by a person skilled in the art.

[29] Approach A: Reduce downstream power based on measured upstream power alone

[30] Downstream power cutback based on received upstream power at the ATU-C is already part of the current standards for the politeness cutback. The politeness cutback is used on very short loops, approximately 0 to 2kft 26AWG equivalent, so there is no impact on downstream capacity regardless of the crosstalk noise environment at the ATU-R. The politeness cutback is performed to minimize the peak signal handling requirements of the ATU-R receiver on short loops. Since it is invoked on only very short loops, it has no impact on power consumption for a typical loop of length 6-8kft.

[31] Referring to FIGS. 1a-1c, it is possible to cutback the downstream power for loops longer than 2kft 26AWG-equivalent and maintain an acceptable transmission rate. However, the degree to which the power is cutback depends on one's willingness to risk a failure during full initialization. For example, a cutback of 12dB is employed where the loop length is estimated to be less than 9kft 26AWG-equivalent. Assuming that the loop length and equivalent downstream loop losses are adequately estimated from the upstream power measured at the ATU-C and there are no adjacent binder T1 disturbers, the 12dB cutback will not significantly reduce the DS capacity below approximately 1.5Mbps.

[32] Unexpected capacity losses may still occur if the ATU-R is not capable of meeting standard test cases because of a noisy front end or amplitude modulation (AM) ingress noise. In those cases, the cutback should be reduced or eliminated and a second full initialization triggered.

[33] Referring to TABLE 1 below, a downstream power cutback is determined as a function of an estimated average of upstream loop attenuation. The estimated average loop attenuation provides an estimate of the loop length. The upstream loop attenuation is defined as the difference between the upstream reference power measured in dBm and the total upstream power measured by the ATU-C on subcarriers 7-18 during stage R-REVERB1 of the initialization. The upstream reference power is defined as the total power used by the ATU-R for transmitting subcarriers 7-18 using the R-REVERB1 transmit power spectral density (PSD) level.

[34]

Average U/S loop attenuation (dB)	>15	>14	>13	>12	>11	>10	>9	>8	>7	>6	>5	>4	>3	>2	>1	>0
D/S Power cut back (dB)	0	1	2	3	4	5	6	6	7	8	8	9	10	11	12	12

[35] TABLE 1

[36] TABLE 1 has been constructed to minimize achievable downstream rate loss with a 6dB SNR margin in the presence of 24 HDSL or 24 DSL near-end crosstalk (NEXT) sources. A person skilled in the art will appreciate that it is also possible to provide different cutback tables based on operator-provisioned parameters. These parameters include a maximum downstream rate and SNR margin. Generally, the cutback is more aggressive when the downstream rate and/or SNR margin are reduced, and less aggressive when they are increased.

[37] The operator can also be given the capability of overriding this mechanism via an additional Operations, Administration and Maintenance (OAM) parameter. It is preferable that the maximum amount of cutback is kept to 12 dB as some implementations may have problems in applying higher cutbacks in the analog domain. A cutback applied in the digital domain results in higher requirements for the dynamic range of a digital-to-analog converter (DAC) in the transmitter.

[38] In order to implement the above in accordance with the current ADSL standards, the ATU-C adjusts its transmit power once the downstream power cutback has been determined. This is achieved early enough (before transmission of C-REVERB) such that the ATU-R can adjust its AGC without triggering a new initialization. Note that this mechanism is similar to the existing politeness cutback but is more aggressive in reducing downstream transmit levels. This cutback is motivated by power savings, unlike the politeness cutback (which attempts to avoid overloading an ATU-R receiver on very short loops).

[39] <u>Approach B: Two passes through full-initialization during initial install or when on-hook loop conditions change significantly</u>

[40] In an alternate embodiment, the ATU-C forces a second full initialization once it has received the B_is, G_is, performance (SNR) margin, and attainable rate

from the ATU-R near the end of the first initialization procedure. A per-carrier SNR (linear scale) at the ATU-R is estimated as $10^{(3*Bi + Gi + margin)/20}$, where i is the carrier number, B_i is the number of bits on carrier i, G_i is the gain of carrier i in dB, and margin is the SNR margin reported by the ATU-R (in dB). The ATU-C uses the downstream SNR to perform a downstream rate (capacity) estimate and determines how much of a cutback can be tolerated for the target downstream rate and margin.

[41] The line rate capacity estimate, less forward error correction (FEC) overhead, in bits per symbol period with no cutback is given by

[42] $$C = \sum_i b_i; \quad b_i = round(\log_2(1 + \frac{SNR_i}{\Gamma})) \quad and\ 2 \le b_i \le b_{\max} \qquad (1)$$

[43] where b_{max} is the maximum number of bits/carrier supported in the ATU-R receiver, and

[44] $\Gamma = 10^{(9.8+10-3)/10}$

[45] is the SNR gap for a bit-error-rate of 10^{-7} with 10dB margin and 3dB coding gain, for G.992.2.

[46] The maximum cutback is determined such that the capacity with that cutback, C', is greater than or equal to $0.98 \times C$, where

[47] $$C' = \sum_i b_i'; \quad b'_i = round(\log_2(1 + \frac{SNR_i / cutback}{\Gamma})) \quad and\ 2 \le b_i \le b_{\max}$$

[48] (2)

[49] The cutback is determined in linear form, corresponding to a multiple of 2dB.

[50] The initialization procedure is then repeated with the ATU-C transmitting at a power level including the cutback throughout the initialization procedure.

[51] It should be noted that the ATU-R B_is, G_is, and performance margin are based on measurements in C-MEDLEY that include both additive noise (e.g., crosstalk) and signal-level-dependent noise, particularly mis-equalization error. The ATU-C has no direct information on the relative contribution of these two noise or error types. If mis-equalization error is the dominant limitation of SNR, then downstream transmit levels could be reduced more aggressively than the case where additive (signal-level-independent) noise is dominant.

[52] This process also requires approximately twice as long (up to about 20 seconds) to initialize the link, but only needs to be invoked whenever a change has occurred

in the noise environment at the ATU-R, for example, when a HDSL service is newly deployed. Such a change in environment may be detected as a degradation in performance during Showtime, including a reduction in the noise margin, an increase in detected CRC errors, and the like. Alternately, an environment change may be detected on a failed power cutback initialization following a period when the communication link is powered down.

[53] Approach C: Reduce DS power by excess margin reported by ATU-R

[54] In another embodiment, negotiation of transmission rates for the communication link occurs over a series of handshakes between the ATU-C and the ATU-R. The ATU-C begins by providing the ATU-R with a list of four (4) rate options to be met. In another alternate embodiment, the ATU-C further transmits the required SNR margin to be met. The ATU-R responds to those options indicating the highest rate, if any, that can be supported. It also provides the ATU-C with the average downstream loop attenuation and the SNR margin at that rate. The values should be approximately the same for all carriers if the G_is have been calculated to equalize the SNR margin across all carriers. If the ATU-R is able to support the highest transmission rate, with a SNR margin in excess of the desired margin, then a power reduction in the amount of the difference between the desired SNR margin and actual SNR margin can be realized. The difficulty with this approach is that the ATU-R is not necessarily designed to adapt to a signal level change at this point in the initialization, forcing a new full initialization.

[55] Approach D: Over-specify margin in first round of rate negotiations

[56] In yet another alternate embodiment, during the first round of rate negotiations, the ATU-C specifies a SNR margin comprising a minimum SNR margin and an additional 'N' dB of SNR margin. If the ATU-R responds (in R-MSG-RA) that it can support a high link rate at the inflated SNR margin, the ATU-C drops its transmitter power by 'N' dB and sets the required margin lower by an equivalent amount for the second round of rate negotiations. If, however, the ATU-R cannot support the target transmission rate with the inflated margin, the ATU-C does not cutback its transmit power, but still reduces the minimum required SNR margin (by 'N') for the second round of negotiations. As with the previous embodiment, the difficulty with this approach is that the ATU-R is not necessarily designed to adapt to a signal level change at this point in the initialization, forcing a new full initialization

[57] More optimal techniques than those described above in Approaches A-D require changes to the existing standards. The changes allow for a faster initialization and maximize the achievable cutback. These techniques are set forth in Approaches E-F below.

[58] <u>Approach E: ATU-R signals attainable cutback implicitly in downstream G_is based on SNR measured in C-MEDLEY</u>

[59] In yet another alternate embodiment the ATU-R can indicate a power cutback implicitly by reducing the G_is to be used on each of the downstream carriers where there is excess SNR margin, based on SNR measurements made during C-MEDLEY. If so, this sets the downstream power cutback desired. With the current ADSL standards, one cannot assume that a vendor's ATU-R will specify the G_is in this manner. That is, the ATU-R may use the G_is only to equalize the margin on each carrier, while keeping an overall unnecessarily large SNR margin. The G_is are communicated to the ATU-C late in the initialization process. As a result, any change in transmit gain must be exact as there is no time for the ATU-R to adapt its receiver to an imprecise gain change before the start of Showtime. Practically, this causes the signal to be reduced in the digital domain, before the digital-to-analog converter (DAC), and can place excess demands on the DAC's dynamic range. This method has the disadvantage of not applying the cutback until the G_i's are implemented on entry to Showtime. As a result, the transmit level, and resulting crosstalk into adjacent lines, remains high through initialization.

[60] As described in Approach B above, the SNR measured by the ATU-R in C-MEDLEY includes both signal-level-dependent and signal-level-independent noise at the ATU-R.

[61] <u>Approach F: ATU-R signals attainable cutback implicitly in downstream G_is based on a signal-level-independent SNR or noise/error signal estimate</u>

[62] In yet another alternate embodiment, the ATU-R determines the amount of downstream power cutback that can be tolerated while meeting the target downstream rate and SNR margin. In order to determine the SNR impact of a reduced downstream transmit signal where mis-equalization error may be dominant, the ATU-R requires several parameters. These parameters include an estimate of the receiver noise floor over the downstream frequency band in the absence of the downstream signal, an estimate of the channel attenuation in the downstream band, and knowledge of the provisioned

downstream rates and margins (i.e., maximums and targets). The ATU-R calculates the achievable downstream cutback from the this information and communicates the cutback back to the ATU-C, which implements the cutback in time for the ATU-R to re-adjust its receiver AGC prior to the SNR estimation during C-MEDLEY. A sample method for estimating the above-described parameters is described as follows.

[63] The ATU-R measures the received downstream power (during C-REVERB1) and estimates the average loop attenuation, as per the current standards, across carriers 7-18 based on a measured per-carrier received signal level and known per-carrier transmit signal levels in C-REVERB1. C-REVERB1 may be sent with a politeness cutback. The ATU-R is informed of the extent of the cutback via a message from the ATU-C. C-REVERB1 may also be replaced by another standard-specific line probing signal that permits extrapolation of the channel attenuation over the downstream frequency band but has different spectral characteristics than those for C-REVERB1.

[64] The ATU-R measures the received downstream noise level during C-QUIET. A minimum noise floor measurement resolution capability may be specified to ensure that the ATU-R is capable of supporting a specific downstream power cutback for a given downstream loop loss and crosstalk noise environment. Any receiver AGC should be set to a value sufficient to resolve the noise level at the loop interface. However, it should not be set so high as to ignore noise contributions from its own front end. The receiver may need to budget for those noise sources if they become dominant at low AGC gain settings.

[65] The ATU-C communicates the maximum, target and minimum downstream transmission rates and SNR margins to the ATU-R. This can be achieved, for example, during C-MSG1 in the current standard.

[66] The ATU-R calculates the amount of downstream cutback possible without reducing estimated downstream capacity below the provisioned maximum downstream rate at the provisioned maximum margin. The capacity calculations are based on a received signal (transmit signal less channel attenuation) and receiver input referenced noise levels as measured at each upstream carrier frequency and mapped to an equivalent per-carrier SNR, SNR_i. Note that this capacity formulation is based on the measured signal and noise levels, independent of near end echo, timing jitter or mis-equalization effects. An additional 6dB of margin is included in the SNR gap figure used in the capacity calculation to cover these impairments, measurement tolerances, and for providing a hedge against introduction of more NEXT sources into the same binder group. The line rate capacity estimates are calculated using equations 1 and 2 as previously described.

[67] The ATU-R communicates the required cutback to the ATU-C, for example in a field of R-MSGS1. Allowing for a power cutback range of 0-15dB in 1dB steps, a 4-bit message field is required.

[68] After receiving and decoding the message from the ATU-R, the ATU-C applies the power cutback approximately 500-3500 symbols into C-MEDLEY. The ATU-R is able to adapt its receiver gain to the new ATU-C transmit level over this interval prior to its final SNR estimation during C-MEDLEY.

[69] FIGS. 2a and 2b are flowcharts generally illustrating the process steps detailed above in Approaches A-F. In FIG. 2a, in step 102 the ATU-C estimates an excess amount of power used for transmitting the downstream signal. The ATU-C may use Approaches A-D to estimate the excess amount of power, as appropriate for each particular implementation. In step 104, the ATU-C reduces the power level of its transmitted (downstream) signal. In step 106, the ATU-C transmits the downstream signal at the reduced power level. The reduced power level achieves a transmission rate of the downstream signal within a predefined tolerance of its target rate.

[70] In FIG. 2b, in step 110 the ATU-R determines an amount of excess power in the downstream signal transmitted from the ATU-C. In step 112, the ATU-R calculates an attainable reduced power level for the downstream signal. In step 114, the ATU-R communicates the reduced power level information to the ATU-C. The ATU-C has enough time to adjust its power level without requiring a second initialization. The ATU-R may use any one or more of Approaches E-F to determine the amount of excess power, calculate the attainable reduced power level, or communicating the reduced power level information, as appropriate.

[71] FIG. 3 is a block diagram of an ATU-C 120 and ATU-R 130 that implement the processes of FIGS. 2a and 2b (as well as Approaches A-F), as appropriate. The ATU-C 120 includes a processor 122 that controls the ATU-C to implement the above-identified processes. The ATU-R 130 includes a processor 132 that controls the ATU-R to implement the above-identified processes. The processors 122 and 132 may be implemented as specific hardware (such as an application-specific integrated circuit or field-programmable gate array) or as general hardware (such as that would implement a computer program or microcode). The ATU-C 120 includes a downstream transmitter 124 and an upstream receiver 126. The ATU-R 130 includes an upstream transmitter 134 and a downstream receiver 136.

[72] Although the invention has been described with reference to certain specific embodiments, various modifications thereof will be apparent to those skilled in the art without departing from the spirit and scope of the invention as outlined in the claims appended hereto.